

Case Study: Improving Financial Projections for Long-Term Care Insurance with Predictive Analytics

By Missy Gordon and Joe Long

Making financial projections is at the heart of what actuaries do. The techniques for doing so have continued to evolve over the years, but the goal is always the same: predict the future as accurately as possible. Nobody can predict the future so there will certainly be fluctuations in financial performance, including the need for additional reserves and capital, but we strive to minimize that fluctuation. In the world of long-term care (LTC) insurance, this is especially challenging for two fundamental reasons: a long projection horizon and complex interactions.

First, the “crystal ball” needs to see 30 years or more into the future as these policies are typically issued to preretirement individuals, but the benefits are often not used for many years into the future. Company data may be limited or nonexistent

at advanced ages, which often requires extrapolation and the need to supplement with industry data.

Additionally, the interactions among variables can be complex, requiring careful construction of the assumption configuration in order to capture the underlying relationships, which can become a daunting task. Company data may also be too limited to capture the true nature of these complex interactions, requiring the use of industry data to understand these relationships.

Traditionally the role of actuarial judgment is often quite large in these efforts to develop projection assumptions, reducing the objectivity and provability of the results. The evolution to using predictive analytics can empower actuaries to quantitatively assess the predictive power of internal versus industry data and determine the “right” balance between the two.


This article is the first of a series that walks through the progression from developing LTC projection assumptions using traditional methods to doing so using predictive analytics. Here we introduce the general concepts. Subsequent articles examine the financial impact of transitioning to predictive analytics in incremental steps, in the context of a case study, for one company where we made this transition.

A BALANCING ACT

In developing a projection assumption, an actuary of even the largest LTC carriers needs to strike a balance between company and industry data.

Trusting the internal data too much may lead to unstable assumptions due to the statistical unreliability of small sample sizes. This is especially true in a business where claims can vary wildly from period to period because of the low frequency and

Figure 1
Illustration of Bias-Variance Trade-off



high severity nature of the claims. However, leaning too heavily on industry data may result in assumptions that are inappropriate for a company's specific blocks of business. In either case, the result is fluctuation against financial projections.

The traditional way to solve this problem is an "actual-to-expected" or "A:E" study. In such a study, experience is compared to an expected assumption (e.g., a benchmark based on industry data), and the actuary applies judgment about data credibility to decide how far from the expected basis to move based on the data. In the traditional approach, balancing the mix of internal and industry data and selecting appropriate variables requires a strong dose of actuarial judgment.

The traditional method has several challenges. First, it is cumbersome. Typically, an actuary uses Microsoft Excel to develop the updated assumption, which can become complex and calculation-intensive. It may be a manual or iterative process, where an expected assumption needs to be updated after determining the A:E adjustment for a given variable before going on to consider an A:E adjustment for another variable. This creates opportunities for human error or assumptions that are not easily reproducible. More importantly, key aspects of the process are judgment-based, including which variables to use, how complex or granular to make the variable interactions, and the amount of weight to give the company's experience. Additionally, the A:E approach typically does not tell us how well the assumptions will work in the future—fit is determined based on the data used to develop the assumption, so a perfect fit does not necessarily mean it will work well for future experience.

When developing a projection assumption it is important for an actuary to give the "right" amount of weight to the experience, while not overreacting to random fluctuations in the data. If one gives too little or too much weight to the data, the assumption will not project future experience well and will lead to financial fluctuations. This is an important concept known as the bias-variance trade-off, which Figure 1 illustrates.

A projection assumption with high bias and low variance tends to be a simple one (e.g., few variables or limited interaction) that gives low weight to the data and typically under-fits the data. Using a single, aggregate A:E adjustment factor may be an example of under-fitting. The projection assumption is highly dependent on the historical mix in the data such that the financial projections will not vary for different mixes of business. The projection assumption may be inappropriate for projecting segments of the business or if the projected mix differs from historical.

On the other hand, a projection assumption that over-fits the data tends to be a complex one (e.g., many variables with granular interactions) that gives high weight to the data and results

in high variance and low bias. Using seriatim A:E adjustment factors is an example of over-fitting. Slight changes in the projected mix will produce wild variations in the financial projections.

The goal is to develop a projection assumption that balances the bias and variance, which the traditional method does using actuarial judgment.

IS THERE A BETTER WAY?

Various predictive analytics techniques can be used to automatically traverse this bias-variance trade-off by determining the "right" amount of data weight that minimizes deviations between future experience and our projections. As our goal is to project future experience as accurately as possible, these techniques provide a robust approach that aligns with our objectives.

... it is critical to question outputs and understand what the model is actually doing.

This in turn reduces the judgment-based decisions relative to how much weight to give the data, which variables to include, and how complex the variable interactions need to be.

One such technique is a penalized generalized linear model (GLM). A GLM develops adjustments to an expected benchmark by giving full weight to the data, but then a penalty is applied to dampen these adjustments. We can think of this penalty as a data weight lever that we use to determine the “right” amount of weight to give the data. A large penalty would give essentially no weight to the company data, leaving the industry benchmark unchanged. On the opposite side, a small penalty gives considerably more weight to the company data and potentially produces large adjustments to the industry benchmark. Using a penalized GLM, the “right” data weight is determined through an automated process that tests a range of data weights and chooses the one that minimizes deviations between unseen experience (data not used in the development of the assumption) and projection assumption.

Using a penalized GLM is a great way to get started with predictive modeling, as it can help us incrementally move from a traditional A:E study to one that uses predictive analytics. We can set up the GLM model in a way such that the only difference from the traditional A:E approach is how much weight is given to the data. The assumptions developed from an A:E and penalized GLM can then be compared side-by-side to get managers, regulators, and auditors comfortable with the new approach.

A penalized GLM approach is very flexible, enabling you to expand and analyze new variables and interactions in the future. Updating a penalized GLM is also simple, and because of the automated process, it is highly repeatable with minimal effort after the initial learning curve and setup. This is in contrast to the cumbersome manual processes often used with traditional A:E methods, which can be slow, prone to human error, and not usually repeatable.

WHAT CHALLENGES REMAIN?

There are challenges that a penalized GLM does not solve, of course.

Actuarial judgment is needed to decide how to extrapolate trends to a future state where there is little to no relevant experience. While the more robust assumptions attributable to penalized GLMs can certainly help in some cases, high levels of variability are to be expected in situations where experience is lacking.

Although industry experience is growing in volume, it can vary wildly across companies because of differences in underwriting, marketing, administration, and plan design. Actuaries working with industry data require great care to ensure they have a solid understanding of the definitions used in the data and their consistency across companies. It is essential that industry data capture key variables to develop a benchmark tailored to a company’s situation. Actuarial judgment is imperative in reviewing the industry data for reasonable relationships before assuming that it is an appropriate expected basis for a company’s situation.

Predictive analytics are powerful tools that require great responsibility. The results must be carefully reviewed to ensure the relationships make sense, for which actuaries are particularly well suited. There can be a temptation to treat any automated process as a black box and simply accept the results, but it is critical to question outputs and understand what the model is actually doing.

STEPPING STONE TO FURTHER EVOLUTION

Once a company gets comfortable with penalized GLMs, it can lead into more powerful machine learning techniques (such as tree-based algorithms) to navigate complex interactions and understand which variables are most important. As a powerful, simple, and well-understood technique, penalized GLMs are a great first foray into the world of predictive modeling.

In our next article, we will dive into a case study and share the results and our experiences in making this transition to predictive analytics to develop LTC claim termination projection assumptions. ■



Missy Gordon, FSA, MAAA, is a principal and consulting actuary at Milliman. She can be reached at missy.gordon@milliman.com.



Joe Long, is an assistant actuary and data scientist at Milliman. He can be reached at joe.long@milliman.com.